

# PREDICTION OF PARTICLE TYPE GIVEN MEASUREMENTS OF PARTICLE LOCATION

BY ROBERT W. JOHNSON

*Alphawave Research*

The MaxEnt approach to the prediction of particle type given measurements of particle location is explored. Two types of particle are considered, and locations are expressed in terms of a single spatial coordinate. Several cases corresponding to different states of prior knowledge are evaluated. Predictions are calculated using the expectation value, which is the average of the observable weighted by the evidence measure over the parameter manifold. How the methodology extends to more general situations is discussed.

**1. Introduction.** A common problem that appears in many guises is how to predict the type of some unidentified particle knowing only its location, given a list of locations for where such particles of identified type have been found. One example of such a problem is the prediction of whether a passerby crossing some particular line in space is male or female based upon knowledge of where on the line previous passersby of known gender have crossed. The problem is predicated on the assumption that measurements of location are inexpensive compared to measurements of classification, so that after an initial training set of data has been evaluated some algorithm can then be used to predict the classification based only upon a known location with some degree of certainty that can be determined.

A specific case of this problem has recently been addressed by [Hall, Park and Samworth \(2008\)](#) using the method of  $k$ th nearest neighbor classification. In that paper, an optimal choice of  $k$  is evaluated empirically from the data, allowing for a nonparametric estimate of the desired probability. In this paper, we will examine how a parametric model is integrated over the parameter manifold with respect to the evidence measure to yield a prediction which depends only upon the given location. The MaxEnt aspect of the algorithm refers to the use of uninformative priors for the parameters, which nonetheless may be nonuniform for some particular choice of coordinate mapping of the parameter manifold.

The process of inductive reasoning is best described using the language of conditional probability theory, as demonstrated by [Bretthorst \(1988\)](#),

---

*AMS 2000 subject classifications:* Primary 62F15, 62H30; secondary 62-07

*Keywords and phrases:* Bayesian inference, inductive reasoning, machine learning

Durrett (1994), and Sivia (1996). Let us quickly review the notation and nomenclature that will be used in this paper. The formal statement of the expression for the probability of  $A$  given conditions  $B$  can be written as

$$(1) \quad p(A | B) \equiv p_B^A ,$$

where  $A$  and  $B$  can have arbitrary dimensionality; for example,  $A$  could be a vector of measurements, and  $B$  could include both the vector of parameters associated with some model as well as any other conditioning statements such as the model index. The sum and product rules of probability theory yield the expressions for marginalization and Bayes' theorem,

$$(2) \quad p^A = \int_{\{B\}} p^{A,B} dB ,$$

$$(3) \quad p_A^B p^A = p_B^A p^B ,$$

where marginalization follows from the requirement of unit normalization, and Bayes' theorem follows from requiring logical consistency of the joint density  $p^{A,B} = p^{B,A}$ . Certain names have come to be associated with the various factors above, but as Sivia (1996) points out, what one calls a probability is irrelevant, as the distinction between  $p_A^C$  and  $p_B^C$  is carried explicitly by the differences in the conditioning statements. Nonetheless, the terms “likelihood” and “prior” are useful for describing how new data updates one's state of knowledge. Instead of the term “posterior” for the estimate of the parameter probability we will use “evidence”, and the chance of measuring the data based on no other knowledge will not be named as it is not necessary for the normalization of the evidence measure nor for the evaluation of the relative evidence for competing models.

**2. Definition of the model.** For consistency of comparison, we will follow as closely as possible the notation used by Hall, Park and Samworth (2008). The population of particles decomposes into two classifications, denoted type  $X$  and type  $Y$ , and the location for each particle type is assumed to follow an independent normal distribution in the spatial dimension  $z$ . The measured locations of the particles identified as type  $X$  can be expressed as the vector  $\mathbf{X} \equiv X_j$  for integer  $j \in [1, J]$ , and similarly  $\mathbf{Y} \equiv Y_k$  for  $k \in [1, K]$ , such that  $N = J + K$  gives the total number of classified particles. The location measurements for type  $X$  are assumed to be drawn from the normal distribution  $p(X_j | \bar{f}, \tilde{f}) \equiv f(X_j)$ , where

$$(4) \quad f(z) = (2\pi\tilde{f}^2)^{-1/2} \exp^{-1/2}[(\bar{f} - z)/\tilde{f}]^2 ,$$

and similarly for  $p(Y_k | \bar{g}, \tilde{g}) \equiv g(Y_k)$ , using notation  $\exp^\alpha(\beta) \equiv e^{\alpha\beta}$  with an economy of brackets when possible. One may alternately interpret that equation to mean that the standard deviation of the location measurement process is equal to  $\tilde{f}$  (or  $\tilde{g}$ ). The relative likelihood of a particle being a  $Y$  rather than an  $X$  is denoted by the parameter  $m$ , such that the absolute likelihood of being an  $X$  is  $p_m^X = (1 + m)^{-1}$ . The parametric model, then, for the probability that some new, unclassified datum is of type  $X$  knowing only its location  $z$  is

$$(5) \quad p(X \text{ at } z | z, m, \bar{f}, \bar{g}, \tilde{f}, \tilde{g}) = [1 + m \zeta(z)]^{-1} ,$$

where  $\zeta(z) \equiv g(z)/f(z)$ , and by normalization  $p_{zm\bar{f}\bar{g}\tilde{f}\tilde{g}}^{Y \text{ at } z} = 1 - p_{zm\bar{f}\bar{g}\tilde{f}\tilde{g}}^{X \text{ at } z}$ . The desired quantity, however, is the estimate of that likelihood given the data  $p_{zNJK\mathbf{XY}}^{X \text{ at } z}$ . Inductive reasoning is used to relate these quantities of interest.

For brevity of notation, any knowledge that can be derived from the conditioning statements explicitly present will be suppressed, *e.g.*  $p_{zNJK\mathbf{XY}}^{X \text{ at } z} = p_{z\mathbf{XY}}^{X \text{ at } z}$ . Let us also collect the coordinates of the parameter manifold into the contravariant position vector  $\mathbf{r} \equiv [m, \bar{f}, \bar{g}, \tilde{f}, \tilde{g}]$ , such that  $\nabla \equiv \partial/\partial\mathbf{r}$  is a covariant vector. By marginalization, the desired quantity can be written as the expectation value of the observable as a function of the parameters weighted by the evidence measure integrated over the entire parameter manifold, not simply taken from its mode (most likely value), formally expressed

$$(6) \quad p_{z\mathbf{XY}}^{X \text{ at } z} = \int_{\{\mathbf{r}\}} p_{z\mathbf{XY}}^{X \text{ at } z, \mathbf{r}} d\mathbf{r} = \int_{\{\mathbf{r}\}} p_{z\mathbf{r}}^{X \text{ at } z} p_{\mathbf{XY}}^{\mathbf{r}} d\mathbf{r} \equiv \langle p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r} | \mathbf{XY}} ,$$

where the conditioning on  $N$  is implicit.

The evidence measure decomposes into a product of factors according to what knowledge is required for their determination  $p_{\mathbf{XY}}^{\mathbf{r}} \propto p_{JK}^m p_{\mathbf{X}}^{\bar{f}\tilde{f}} p_{\mathbf{Y}}^{\bar{g}\tilde{g}}$ , which themselves factor into products of likelihoods and priors. For example, the evidence density for  $m$  is  $p_{JK}^m \propto p_m^{JK} p_m^m$ , and similarly for the remaining factors. Addressing first the likelihood factors, the chance of observing  $K$  out of  $N$  particles of type  $Y$  given  $m$  is

$$(7) \quad p_m^{JK} = m^K / (1 + m)^{J+K} = m^K (1 + m)^{-N} ,$$

whereas the chance of observing locations  $\mathbf{X}$  given values for  $\bar{f}$  and  $\tilde{f}$  is

$$(8) \quad p_{\bar{f}\tilde{f}}^{\mathbf{X}} = \prod_{z \in \mathbf{X}} f(z) = (2\pi\tilde{f}^2)^{-J/2} \exp^{-1/2} \sum_{z \in \mathbf{X}} [(\bar{f} - z)/\tilde{f}]^2 ,$$

and similarly  $p_{\bar{g}\tilde{g}}^{\mathbf{Y}} = \prod_{z \in \mathbf{Y}} g(z)$ . The likelihood of the data at manifold position  $\mathbf{r}$  is then given by their product  $p_{\mathbf{r}}^{\mathbf{XY}} = p_m^{JK} p_{\bar{f}\tilde{f}}^{\mathbf{X}} p_{\bar{g}\tilde{g}}^{\mathbf{Y}}$ .

Next let us look at the prior factors  $p^{\mathbf{r}} = p^m p^{\bar{f}\tilde{f}} p^{\bar{g}\tilde{g}}$ . The MaxEnt approach to selecting an uninformative prior is based on the principle of indifference as represented by the requirement of consistency under various transformations of the parameter or data coordinate mappings. [Dose \(2003\)](#) gives an excellent description of the process. According to [Jaynes \(1968\)](#), the prior suggested by Jeffreys for the parameters of a Gaussian distribution results from satisfaction of the functional equation for transformations in location and scale, thus  $p^{\bar{f}\tilde{f}\bar{g}\tilde{g}} \propto (\tilde{f}\tilde{g})^{-1}$ . In that paper, he argues that the same functional form is appropriate for the rate parameter of a Poisson process, which will be called  $n$ . In the presentation by [Hall, Park and Samworth \(2008\)](#), the arrival of the particles at the axis of measurement is assumed to follow the Poisson distribution with parameters  $\mu$  and  $\nu$  for types  $X$  and  $Y$  respectively, thus the total number of particles expected in one unit of time is  $n = \mu + \nu$ . Writing  $m = \nu/\mu$ , the determinant of the Jacobian is  $\det J = n/(1+m)^2$ , leading to the transformation of the prior  $p^{nm} = p^{\mu\nu} \det J$ , whereby

$$(9) \quad p^m \propto n^2(1+m)^{-2}\mu^{-1}\nu^{-1} = m^{-1},$$

which states that  $m$  is just as likely to be between 0.1 and 1 as it is to be between 1 and 10 before any observations are recorded. The prior measure for the parameter manifold is thus  $p^{\mathbf{r}} \propto (m\tilde{f}\tilde{g})^{-1}$ , where the constant of proportionality is determined by the limits of consideration. In particular, finite limits for  $m$  must be symmetric in scale about unity so that the prior expectation of finding an  $X$ ,

$$(10) \quad \langle p_m^X \rangle_{m|m_\infty} \equiv \frac{\int_{1/m_\infty}^{m_\infty} m^{-1}(1+m)^{-1} dm}{\int_{1/m_\infty}^{m_\infty} m^{-1} dm},$$

remains equal to 1/2 and so that  $\langle m \rangle_{m|m_\infty} = \langle m^{-1} \rangle_{m|m_\infty}$ .

**3. Comparison to simpler models.** If the parameters  $\bar{f}$ ,  $\bar{g}$ ,  $\tilde{f}$ , and  $\tilde{g}$  for the Gaussian distributions are known in advance, then only  $m$  need be estimated from the data. Using the notation  $q_\beta^\alpha \equiv -\log p_\beta^\alpha$ , the “parameter info” (information content of the evidence density) is

$$(11) \quad q_{JK}^m = q_m^{JK} + q^m + C = N \log(1+m) - (K-1) \log m + C,$$

where  $C$  is the logarithm of the normalizing constant, equal to  $\log \beta(J, K)$  when  $m_\infty = \infty$ . Its mode may be found from the equation for a vanishing gradient

$$(12) \quad 0 = \partial_m q_{JK}^m = m^{-1}(1+m)^{-1}[1 - K + m(J+1)],$$




FIG 1. Comparison of  $\text{prob}(X \text{ at } z)$  as a function of  $\zeta(z)$  estimated from the expectation value  $\langle p_{z\zeta_m}^{X \text{ at } z} \rangle_{m|JK}$  shown as  $\square$  to its approximation  $p_{z\zeta\langle m \rangle}^{X \text{ at } z}$  shown as  $\circ$  for  $J$  and  $K$  as indicated above each panel.

whose solution  $m_0 = (K-1)/(J+1)$  may be negative when  $K=0$ , in which case the mode is at the lower limit of consideration; if one inverts the identity of  $X$  and  $Y$ , one finds that the evidence for the inverted  $m$  has its mode at the position given by the analytic formula. In the limit of  $m \in [0, \infty]$ , the expectation that some new particle is of type  $X$  is  $\langle p_m^X \rangle_{m|JK} = (1+K/J)^{-1}$ , the expectation value for  $m$  is  $\langle m \rangle_{m|JK} = K/(J-1)$ , and that for  $m^{-1}$  is  $\langle m^{-1} \rangle_{m|JK} = J/(K-1)$ . When the new particle's location  $z$  also is known, one may evaluate  $\langle p_{z\zeta_m}^{X \text{ at } z} \rangle_{m|JK}$  as the prediction for its type being  $X$ ; when both  $J$  and  $K$  are large, that value is approximately given by simply using the expectation for  $m$  in the model for  $p_{z\zeta_m}^{X \text{ at } z}$ , *i.e.*  $\langle p_{z\zeta_m}^{X \text{ at } z} \rangle_{m|JK} \approx p_{z\zeta\langle m \rangle}^{X \text{ at } z}$  for  $J, K \gg 1$ . In Figure 1 we compare  $\langle p_{z\zeta_m}^{X \text{ at } z} \rangle_{m|JK}$  to  $p_{z\zeta\langle m \rangle}^{X \text{ at } z}$  as a function of  $\zeta(z)$  for various values of  $J$  and  $K$ .

Let us next consider the case where only the deviations  $\tilde{f}$  and  $\tilde{g}$  are known in advance, so that now  $\bar{f}$  and  $\bar{g}$  must also be estimated from the data. For convenience, let us further suppose that  $\tilde{f} = \tilde{g} = 1$ , setting the unit for the locations  $z$ . The values of the parameters used to generate the data will be denoted  $m_\Omega$ ,  $\bar{f}_\Omega$ , and  $\bar{g}_\Omega$ , where the subscript  $\Omega$  indicates conditioning on the sum of all knowledge, *i.e.* they are the “true” values unknown to mere mortals. According to the model, the actual chance of finding an  $X$  at some  $z$  is given by  $p_{z\Omega}^{X \text{ at } z} = [1 + m_\Omega \zeta_\Omega(z)]^{-1}$ , thus the estimation of that quantity, as well as its reliability, from the data at hand is the desired goal of the

statistical analysis. The manifold position  $\mathbf{r}_\Omega$  of course is not allowed to be part of that process, as its knowledge would preclude the need to collect any data.

Retaining only  $m$ ,  $\bar{f}$ , and  $\bar{g}$  in  $\mathbf{r}$ , the evidence density is

$$(13) \quad p_{\mathbf{XY}}^{\mathbf{r}} \propto p_m^{JK} p_{\bar{f}}^{\mathbf{X}} p_{\bar{g}}^{\mathbf{Y}} p^{m\bar{f}\bar{g}},$$

where the prior is normalized to unit volume, thus the parameter info is

$$(14) \quad q_{\mathbf{XY}}^{\mathbf{r}} = q_m^{JK} + q^m + q_{\bar{f}}^{\mathbf{X}} + q_{\bar{g}}^{\mathbf{Y}} + C,$$

and the value of  $C$  is chosen according to the task at hand. For taking expectation values, one convenient choice is that which normalizes the peak of the evidence to unity  $C_0 = -q_0$ , whereas for hypothesis testing (model selection)  $C$  must equal the logarithm of the normalizing constant for the prior, and the remaining terms must retain any constants found in the likelihood, unless they happen to cancel out of the relative evidence ratio. Here,  $C_{\mathbf{r}} = \log(2 \log m_\infty) + 2 \log \Delta_z$  when  $m \in [m_\infty^{-1}, m_\infty]$  and  $\bar{f}, \bar{g} \in [-\Delta_z/2, \Delta_z/2]$ . The expression

$$(15) \quad q_{\bar{f}}^{\mathbf{X}} + q_{\bar{g}}^{\mathbf{Y}} = \frac{N}{2} \log(2\pi) + \frac{1}{2} \sum_{z \in \mathbf{X}} (\bar{f} - z)^2 + \frac{1}{2} \sum_{z \in \mathbf{Y}} (\bar{g} - z)^2$$

gives the additional likelihood info coming from the location measurements, and its mode is easily found to be at  $\bar{f}_0 = \langle X_j \rangle_j$  and  $\bar{g}_0 = \langle Y_k \rangle_k$ .

Because the preferred locations  $\bar{f}$  and  $\bar{g}$  are themselves estimated, one must ask whether the data would be more efficiently represented by a single Gaussian, which we will call  $h(z)$  with mean  $\bar{h}$ . The relevant factors in the relative evidence ratio are those which do not depend on  $m$ , and the problem reduces to the well-known example of whether the difference in means between two populations is statistically significant. The answer is given by the ratio of the expectation values for the likelihood of each model, (16)

$$\rho_h^{fg} \equiv \frac{\langle p_{\bar{f}\bar{g}}^{\mathbf{XY}} \rangle_{\bar{f}\bar{g}}}{\langle p_{\bar{h}}^{\mathbf{XY}} \rangle_{\bar{h}}} = \frac{\int_{-\Delta_z/2}^{\Delta_z/2} \int_{-\Delta_z/2}^{\Delta_z/2} p_{\bar{f}\bar{g}}^{\mathbf{XY}} p^{\bar{f}\bar{g}} d\bar{f} d\bar{g}}{\int_{-\Delta_z/2}^{\Delta_z/2} p_{\bar{h}}^{\mathbf{XY}} p^{\bar{h}} d\bar{h}} \approx \left( \frac{p_{\bar{f}_0}^{\mathbf{X}} p_{\bar{g}_0}^{\mathbf{Y}}}{p_{\bar{h}_0}^{\mathbf{XY}}} \right) \left( \frac{2\pi N}{\Delta_z^2 JK} \right)^{1/2},$$

where the first factor is the ratio of peak likelihoods (when the prior is uniform), the second (Occam) factor is the ratio of the filling fractions for each model, and the approximation results from taking infinite bounds in the integrals. The filling fraction for a model is a number between 0 and 1 which indicates how much the evidence fills the parameter manifold with

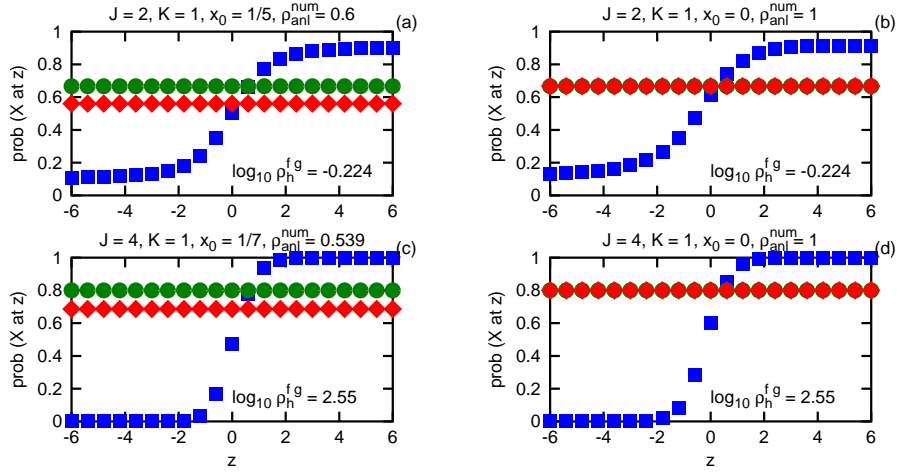


FIG 2. Comparison of the numerical estimate for  $\langle p_m^X \rangle_{m|JK}$  displayed as  $\diamond$  to that for  $\langle p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\mathbf{XY}}$  displayed as  $\square$  using limits of integration given by  $\Delta_z = 12$  and  $x_0$  as indicated above each panel, with the analytic value for  $\langle p_m^X \rangle_{m|JK}$  displayed as  $\circ$ .

respect to the prior measure. If the relative evidence ratio is well below unity  $\rho_h^{fg} \ll 1$ , then probability theory is telling one to neglect the  $z$  dependence and simply use  $m$  as the basis for any prediction regarding the type of some unidentified particle. If either  $J$  or  $K$  equals 0, attention to factors reveals that  $\rho_h^{fg} = 1$  in that case, as the evidence density is uniform along the irrelevant parameter.

Using finite limits for numerical integration can have an impact on the result. Recalling that conditioning on  $N$  has been implied throughout, let us compare the estimate of  $\langle p_m^X \rangle_{m|JK}$  to  $\langle p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\mathbf{XY}}$  for the two cases of  $m_\infty = N+1$  and  $m_\infty = \infty$ . The former represents a prior state of knowledge in which one is certain to observe particles of both types eventually, even though particles of only one type have been observed so far. The numerical integration is more easily accomplished upon a change of variables  $x = (1+m)^{-1}$  such that

$$(17) \quad \int_{1/m_\infty}^{m_\infty} m^{K-1} (1+m)^{-J-K} dm = \int_{x_0}^{1-x_0} x^{J-1} (1-x)^{K-1} dx \leq \beta(J, K),$$

with equality in the limit  $x_0 \rightarrow 0$ . One way to assess the reliability of one's estimate is to inspect the ratio of the numerically integrated evidence volume to that derived analytically for the infinite manifold, identified as  $\rho_{\text{anl}}^{\text{num}}$ . Figure 2 shows the comparison of the estimated operators for values of  $N = 3$  and  $N = 5$  using  $\bar{f}_\Omega = 1$  and  $\bar{g}_\Omega = -1$ , as well as the analytic expression for

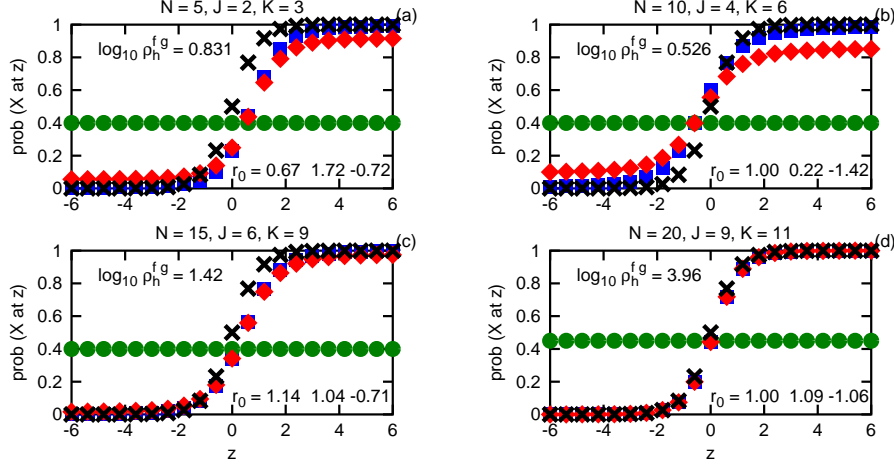


FIG 3. Comparison of the analytic value for  $\langle p_m^X \rangle_{m|JK}$  displayed as  $\circ$  to the numerical estimate for  $\langle p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\mathbf{XY}}$  displayed as  $\square$  for values of  $N$  as indicated above each panel, with their weighted mean  $p_{z\rho}^{X \text{ at } z}$  displayed as  $\diamond$  and the underlying distribution  $p_{z\Omega}^{X \text{ at } z}$  displayed as  $\times$ . Also shown are the values for the parameter mode  $\mathbf{r}_0$ .

$\langle p_m^X \rangle_{m|JK}$  when  $m_\infty = \infty$ . With only a few measurements, the numerical estimate for  $\langle p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\mathbf{XY}}$  approaches the analytic value for  $\langle p_m^X \rangle_{m|JK}$  only when  $m_\infty$  is very large. After more measurements have accumulated, the limits on  $m$  have less impact; however, if the mode in  $m$  is close to the numerical limit, the estimate may still be inaccurate.

To collect the estimates from the two possible models into a single prediction, one simply averages them with weights given by their relative evidence,

$$(18) \quad p_{z\rho}^{X \text{ at } z} \equiv (\rho_h^{fg} \langle p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\mathbf{XY}} + \langle p_m^X \rangle_{m|JK}) / (\rho_h^{fg} + 1) .$$

As evidence accumulates in favor of the  $f/g$  model, their average quickly approaches the estimate from just that model, as seen in Figure 3. Using values of  $m_\Omega = 1$ ,  $\bar{f}_\Omega = -\bar{g}_\Omega = 1$  and integration limits of  $\Delta_z = 12$  and  $x_0 = 0$ , it takes only a few tens of measurements before the evidence for the null hypothesis (the  $h$  model) is negligible. As the number of measurements in the training data grows, the estimate  $p_{z\rho}^{X \text{ at } z}$  converges to the underlying  $\Omega$  distribution for the chance of finding an  $X$  at  $z$ .

So far we have said nothing about the rate of convergence of the estimate  $\langle p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\mathbf{XY}}$  as a function of the number of measurements  $N$ . Partly that is because we have summarized our inference about  $p_{z\mathbf{XY}}^{X \text{ at } z}$  into a single number, the expectation value of the observable  $p_{z\mathbf{r}}^{X \text{ at } z}$ , and to answer the question of convergence requires keeping track of two numbers for the inference, rep-



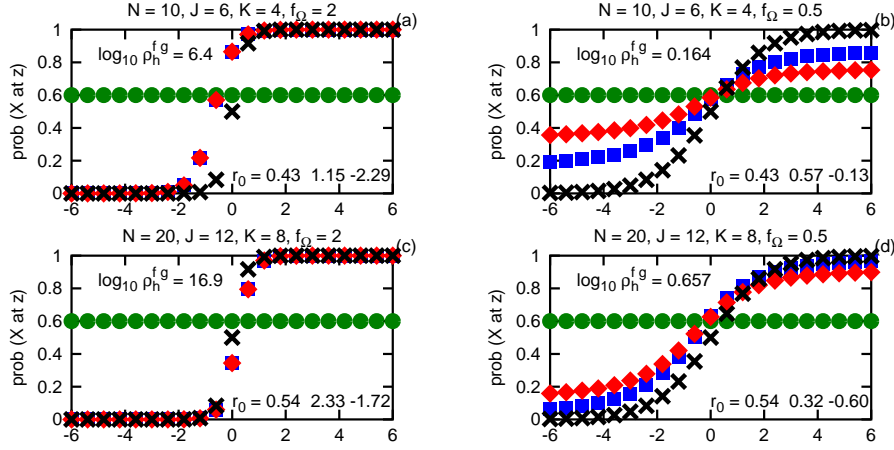


FIG 4. Comparison of the analytic value for  $\langle p_m^X \rangle_m |_{JK}$  displayed as  $\circ$  to the numerical estimate for  $\langle p_{zr}^{X \text{ at } z} \rangle_r |_{\mathbf{XY}}$  displayed as  $\square$  for values of  $N$  and  $\bar{f}_\Omega = -\bar{g}_\Omega$  as indicated above each panel, with their weighted mean  $p_{z\rho}^{X \text{ at } z}$  displayed as  $\diamond$  and the underlying distribution  $p_{z\Omega}^{X \text{ at } z}$  displayed as  $\times$ . Also shown are the values for the parameter mode  $\mathbf{r}_0$ .

representing for example its central location and its width. That procedure will be addressed later in this article. While we have blithely displayed  $p_{z\Omega}^{X \text{ at } z}$  in the preceding figure, one should never forget that its knowledge is beyond the ken of mortals. Practically speaking, one simply must collect a sufficient amount of data such that collecting more data no longer significantly influences the estimate, implicitly assuming that the underlying physical process is stationary in time.

The ratio of the the deviation in the data  $\tilde{f}$  to the separation of the preferred locations  $\tilde{f} - \bar{g}$  affects how much data is necessary for convergence of the estimate. When that ratio is small, not many measurements are needed before the null hypothesis is discounted, and further measurements serve only to improve the convergence. However, when that ratio is large, the null hypothesis can be discounted only after a sufficient number of measurements have been taken so that the parameter evidence is well resolved. In Figure 4 we show the same estimates as in Figure 3 but for  $N$  equal to 10 and 20 and  $\bar{f}_\Omega = -\bar{g}_\Omega$  of 2 and 1/2. Even when the mode  $\mathbf{r}_0$  gives a poor reckoning of the underlying process, the expectation value  $\langle p_{zr}^{X \text{ at } z} \rangle_r |_{\mathbf{XY}}$  is fairly accurate when  $\rho_h^{fg} \gg 1$ . To make from  $\rho_h^{fg}$  a number comparable to the  $P$  (or  $Q$ ) value of frequentist methods, one would state that the null hypothesis ( $h$  model) is discounted at the level of  $(\rho_h^{fg} + 1)^{-1}$ .

**4. Integration over unknown deviation.** Let us now consider the more realistic situation where the deviations of the location distributions are not known in advance. For convenience, let us assign them all the domain of  $\sigma \in [\sigma_0, \sigma_\infty]$  with prior  $p^\sigma \propto \sigma^{-1}$ . The normalization constant is  $\Delta_{\log \sigma} \equiv \log \sigma_\infty - \log \sigma_0$ , which equals infinity if either  $\sigma_\infty = \infty$  or  $\sigma_0 = 0$ . The first task is to evaluate the relative evidence of the models,

$$(19) \quad \rho_h^{fg} = \frac{\int_{-\Delta_z/2}^{\Delta_z/2} \int_{-\Delta_z/2}^{\Delta_z/2} \int_{\sigma_0}^{\sigma_\infty} \int_{\sigma_0}^{\sigma_\infty} \tilde{f}^{-J-1} \tilde{g}^{-K-1} \exp^{-1/2}(\chi_f^2 + \chi_g^2) d\tilde{f} d\tilde{g} d\bar{f} d\bar{g}}{\Delta_z \Delta_{\log \sigma} \int_{-\Delta_z/2}^{\Delta_z/2} \int_{\sigma_0}^{\sigma_\infty} \tilde{h}^{-N-1} \exp^{-1/2}(\chi_h^2) d\tilde{h} d\bar{h}},$$

where  $\chi_f^2 \equiv \sum_{z \in \mathbf{X}} [(\bar{f} - z)/\tilde{f}]^2$  and similarly for  $\chi_g^2$  and  $\chi_h^2$ , which can be approximated as before by taking the Gaussian integrals over infinite limits yet retaining the finite normalization to yield

$$(20) \quad \rho_h^{fg} \approx \left( \frac{2\pi N}{\Delta_z^2 JK} \right)^{1/2} \frac{\int_{\sigma_0}^{\sigma_\infty} \int_{\sigma_0}^{\sigma_\infty} \tilde{f}^{-J} \tilde{g}^{-K} \exp^{-1/2}(\xi_f^2/\tilde{f}^2 + \xi_g^2/\tilde{g}^2) d\tilde{f} d\tilde{g}}{\Delta_{\log \sigma} \int_{\sigma_0}^{\sigma_\infty} \tilde{h}^{-N} \exp^{-1/2}(\xi_h^2/\tilde{h}^2) d\tilde{h}},$$

where  $\xi_f^2 \equiv J(\langle X_j^2 \rangle_j - \langle X_j \rangle_j^2)$  and similarly for  $\xi_g^2$  and  $\xi_h^2$ . The remaining integrals can be evaluated analytically to give the result

$$(21) \quad \rho_h^{fg} \approx \frac{(\pi N / JK)^{1/2} \xi_h^{N-1} \Delta_\Gamma(\xi_f) \Delta_\Gamma(\xi_g)}{2 \Delta_z \Delta_{\log \sigma} \xi_f^{J-1} \xi_g^{K-1} \Delta_\Gamma(\xi_h)},$$

where  $\Delta_\Gamma$  is defined in terms of the upper incomplete gamma function  $\Gamma(a, z) \equiv \int_z^\infty t^{a-1} e^{-t} dt$ , such that

$$(22) \quad \Delta_\Gamma(\xi_f) \equiv \Gamma\left(\frac{J-1}{2}, \frac{\xi_f^2}{2\sigma_\infty^2}\right) - \Gamma\left(\frac{J-1}{2}, \frac{\xi_f^2}{2\sigma_0^2}\right),$$

and similarly for  $\Delta_\Gamma(\xi_g)$  and  $\Delta_\Gamma(\xi_h)$ . Note that this formulation makes no use of the peak evidence ratio but instead is expressed entirely in terms of the data and the limits of the prior.

To recover the peak evidence ratio (and thus the Occam factor), one needs to evaluate the model evidence densities at their mode positions. The parameter info for the  $\mathbf{X}$  data is now

$$(23) \quad q_{\mathbf{X}}^{\tilde{f}\tilde{f}} = (J+1) \log \tilde{f} + (2\tilde{f}^2)^{-1} \sum_{z \in \mathbf{X}} (\bar{f} - z)^2 + C,$$

whose gradient is given by

$$(24) \quad \nabla q_{\mathbf{X}}^{\tilde{f}\tilde{f}} = \tilde{f}^{-3} \left[ \frac{\tilde{f} \sum_{z \in \mathbf{X}} (\bar{f} - z)}{(J+1)\tilde{f}^2 - \sum_{z \in \mathbf{X}} (\bar{f} - z)^2} \right],$$

which vanishes at the mode  $\nabla q_{\mathbf{X}}^{\tilde{f}\tilde{f}}(\bar{f}_0, \tilde{f}_0) = 0$ . As before, the mode in  $\bar{f}$  is at the mean of the locations  $\bar{f}_0 = \langle X_j \rangle_j$ , and the mode in  $\tilde{f}$  can be written as  $\tilde{f}_0 = (J+1)^{-1/2} \xi_f$ . The position of the mode for the remaining Gaussian parameters is found similarly.

Let us briefly discuss the limits of integration hence the normalization of the prior. If the prior is not to be based upon the current crop of measurements, where does the information for the limits come from? The practical answer is that the limits are determined by the nature of the measurement apparatus. Any set of measurements collected within a finite span of time necessarily are limited by the range and resolution of the device used for their collection, for example measurements of the voltage of a circuit collected by a common voltmeter. As long as no measurement “pegs the needle” one can safely use limits based on the range of the device; those that do can be addressed through an appropriately modified contribution to the likelihood beyond the scope of this article. Similarly, the resolution of the device (or the width of the particle) sets a lower limit on what can be said about any measured deviations in the population locations. For the evaluation of the evidence ratio  $\rho_h^{fg}$  as well as  $\rho_{\text{anl}}^{\text{num}}$ , we will set the limits for the deviations as  $\sigma_\infty = \Delta_z$  and  $\sigma_0 = 10^{-4} \Delta_z$ , with  $\Delta_z = 12$  and  $x_0 = 0$  as above.

The evaluation of the expectation value of the observable  $\langle p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\mathbf{X}\mathbf{Y}}$  proceeds as before, only now the integration is over a 5 dimensional parameter manifold. As [Press et al. \(1992\)](#) state, “integrals of functions of several variables, over regions with dimension greater than one, are *not easy*.” For problems of Bayesian inference, the majority of the contribution to the integral comes from a region localized around the peak of the evidence density when sufficient data exists that the limits of the prior are irrelevant. Luckily, for this problem the evidence mode is unique and analytic, so that one may select limits for the numerical integration much tighter than those given by the prior while still encompassing 99.9% of the normalized evidence density. The evaluation is performed using an adaptive grid algorithm ([Genz and Malik, 1980](#); [Berntsen, Espelid and Genz, 1991](#)) over a small region of the manifold centered on the position of the optimal parameter values.

In [Figure 5](#) we compare the expectation of the observable  $\langle p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\mathbf{X}\mathbf{Y}}$  to that given by evaluating the observable using the parameter mode  $p_{z\mathbf{r}_0}^{X \text{ at } z}|\mathbf{X}\mathbf{Y}$ , with the results given in [Table 1](#). As the number of measurements increases, those estimates draw closer, according to the narrowing of the peak in the evidence density. The estimate from the expectation value is “more conservative” than that from the mode, in that it is closer to the estimate  $\langle p_m^X \rangle_m|\mathbf{X}\mathbf{Y}$  (not shown). While the estimate from the mode more closely resembles the underlying distribution  $p_{z\Omega}^{X \text{ at } z}$ , in a real world situation we are not privy

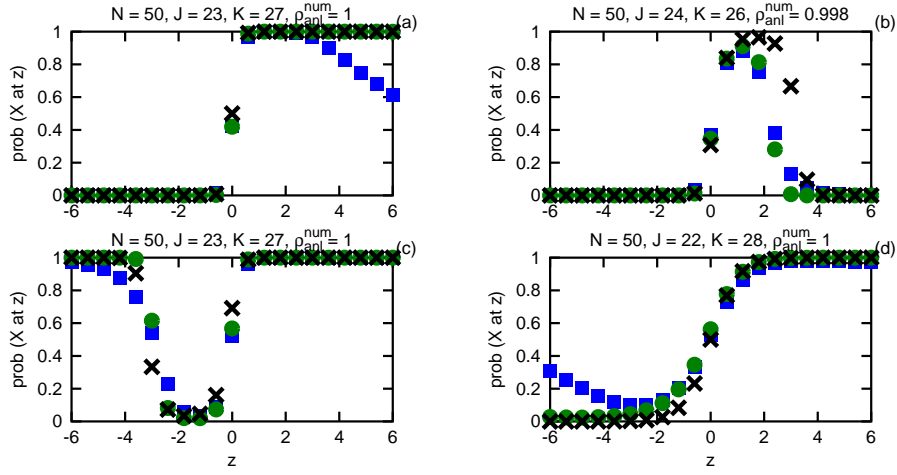


FIG 5. Comparison of the expectation value  $\langle p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\mathbf{XY}}$  displayed as  $\square$  to the estimate from the parameter mode  $p_{z\mathbf{r}_0}^{X \text{ at } z}|\mathbf{XY}$  displayed as  $\circ$  for  $N = 50$  as well as the underlying distribution  $p_{z\Omega}^{X \text{ at } z}$  displayed as  $\times$  with parameter values as given in Table 1.

to that knowledge (which would obviate the need for a statistical analysis). The expectation value  $\langle p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\mathbf{XY}}$  summarizes what the data have to say about the observable into a single number. Using  $p_{z\mathbf{r}_0}^{X \text{ at } z}|\mathbf{XY}$  as a proxy for the mode of the observable (which is not quite the same thing as the observable of the mode), one could ascribe to  $p_{z\mathbf{r}}^{X \text{ at } z}$  a beta distribution according to its mean and mode; the maximum likelihood estimate of the parameters of the beta distribution for the observable  $p_{z\mathbf{r}}^{X \text{ at } z}$  requires evaluation of the two observables  $\langle \log p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\mathbf{XY}}$  and  $\langle \log(1 - p_{z\mathbf{r}}^{X \text{ at } z}) \rangle_{\mathbf{r}|\mathbf{XY}}$ .

TABLE 1  
Parameters corresponding to Figure 5

panel	$\log_{10} \rho_h^{fg}$	$N$	$m_\Omega$	$\bar{f}_\Omega$	$\bar{g}_\Omega$	$\tilde{f}_\Omega$	$\tilde{g}_\Omega$
	$J$	$K$	$m_0$	$f_0$	$g_0$	$\tilde{f}_0$	$\tilde{g}_0$
a	13.8	50	1	1	-1	0.5	0.5
	23	27	1.08	0.879	-1.01	0.415	0.541
b	8.31	50	1	1	-1	0.5	1
	24	26	1	0.844	-0.837	0.44	1.09
c	12.1	50	1	1	-1	1	0.5
	23	27	1.08	1.22	-0.875	0.933	0.443
d	2.24	50	1	1	-1	1	1
	22	28	1.17	0.446	-1.08	1.11	0.945

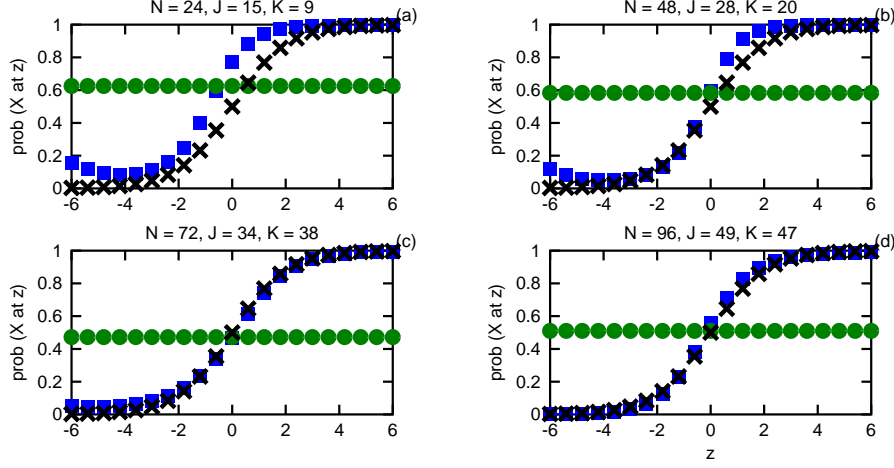


FIG 6. Comparison of the expectation value  $\langle p_{zr}^{X \text{ at } z} \rangle_{\mathbf{r}} | \mathbf{X} \mathbf{Y}$  displayed as  $\square$  to  $\langle p_m^X \rangle_m | \mathbf{X} \mathbf{Y}$  displayed as  $\circ$  for various  $N$  as well as the underlying distribution  $p_{z\Omega}^{X \text{ at } z}$  displayed as  $\times$  with parameter values as given in Table 2.

**5. Comparison to non-parametric classification.** Let us now look at how the MaxEnt prediction compares to those derived from some non-parametric schemes commonly employed for this type of problem. For this section we will use the same data for each method generated using parameter values  $m_\Omega = 1$ ,  $\bar{f}_\Omega = -\bar{g}_\Omega = 2$ , and  $\tilde{f}_\Omega = \tilde{g}_\Omega = 2$  for various  $N$ . We will consider both a nearest neighbor classification scheme which produces its estimate from a subset of the data “close” to the desired location as well as a classification scheme based upon a kernel density estimate of the identified particle distributions. The MaxEnt estimates from these sets of data are shown in Figure 6 with parameter modes in Table 2.

Starting with nearest neighbor classification, its prediction for the type of some new particle,

$$(25) \quad p_{z\kappa\mathbf{X}\mathbf{Y}}^{X \text{ at } z} \equiv \kappa_X / (\kappa_X + \kappa_Y),$$

is conditioned on the number of neighbors  $\kappa = \kappa_X + \kappa_Y$  chosen to be influential, denoted here as  $\kappa \equiv \rho_\kappa N$  for  $0 < \rho_\kappa \leq 1$  such that  $\kappa$  is an integer. The selection of  $\kappa$  is arbitrary, but Hall, Park and Samworth (2008) describe a method for choosing its value based upon bootstrap estimates from the data. Here, however, we are interested in the case where there is so little data that bootstrap estimates are unlikely to be reliable. Consequently, we will consider the set of ratios  $\rho_\kappa \in [1/8, 1/4, 1/2, 1]$ . Another distinction is that they assign a type of  $X$  or  $Y$  to the new particle at  $z$  according to

TABLE 2  
Parameters corresponding to Figure 6

panel	$\log_{10} \rho_h^{fg}$	$N$	$m_\Omega$	$\bar{f}_\Omega$	$\bar{g}_\Omega$	$\bar{f}_\Omega$	$\bar{g}_\Omega$
	$J$	$K$	$m_0$	$f_0$	$g_0$	$\tilde{f}_0$	$\tilde{g}_0$
a	3.09	24	1	2	-2	2	2
	15	9	0.5	2.02	-2.58	2.05	1.22
b	7.82	48	1	2	-2	2	2
	28	20	0.655	2.3	-1.99	1.98	1.27
c	6.77	72	1	2	-2	2	2
	34	38	1.06	1.58	-1.62	1.92	1.64
d	15.8	96	1	2	-2	2	2
	49	47	0.92	2.08	-2.52	1.76	2.11

whether  $p_{z\kappa\mathbf{X}\mathbf{Y}}^{X \text{ at } z}$  is greater or less than  $1/2$  rather than expressing the chance of finding an  $X$  at  $z$  as a probability.

The results of the preceding method are shown in Figure 7. One feature of the nearest neighbors method is that its prediction is quantized in units of  $1/\kappa$ , which can lead to large jumps in the estimate when there is not much data. These jumps yield an estimate which is not smooth as a function of  $z$ , even as the number of measurements approaches 100. By basing its prediction on the rank of the distances from the data to the desired location, this method throws away information pertinent to the analysis. Consequently, its prediction is a coarsely grained representation of the underlying distribution, even for moderately large sets of data.

Alternately, we can consider a classification scheme based upon a kernel density estimate of the particle distributions by type. Its prediction  $p_{z\lambda\mathbf{X}\mathbf{Y}}^{X \text{ at } z}$  is conditioned on the bandwidth parameter  $\lambda$  for the kernel resolution. The kernel basis chosen is that given by the Gaussian distributions  $\delta_\lambda(z) \equiv \exp^{-1/2}(z^2/\lambda^2)$ , with peaks normalized to unity for convenience later. The kernel density estimate for the distribution of type  $X$  is equal to the sum of the kernel basis functions centered on the datum locations  $f_{\lambda\mathbf{X}}(z) \equiv \sum_j \delta_\lambda(z - X_j)$ , and similarly for  $g_{\lambda\mathbf{Y}}(z)$ . The classification prediction is then determined from the ratio

$$(26) \quad p_{z\lambda\mathbf{X}\mathbf{Y}}^{X \text{ at } z} \equiv f_{\lambda\mathbf{X}}(z)/[f_{\lambda\mathbf{X}}(z) + g_{\lambda\mathbf{Y}}(z)] ,$$

using bandwidth parameters of  $\lambda \in [1/4, 1, 4, 16]$ .

The results for the preceding method are shown in Figure 8. Again, large jumps are apparent in the estimate when the kernel width is much smaller than what we know the particle deviations to be, but the values are continuous rather than being quantized. The effect of the kernel basis is that of a smoothing filter which spreads the information from each datum over

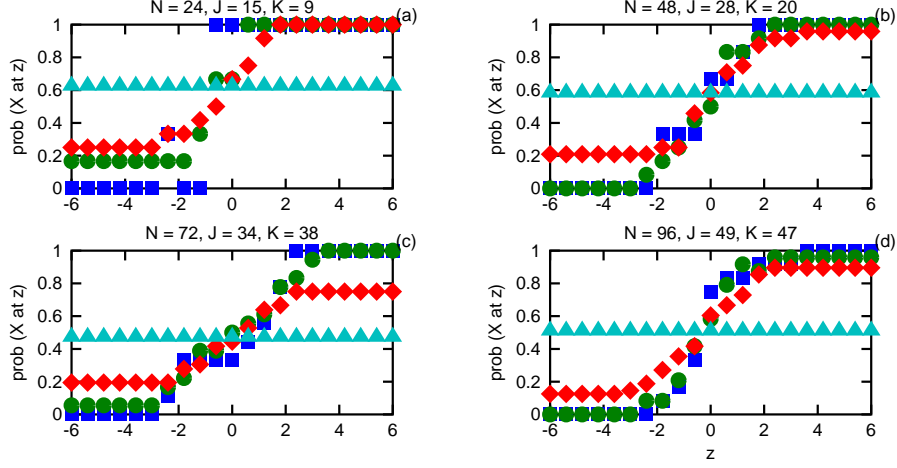


FIG 7. Comparison of the nearest neighbor estimates  $p_{z\kappa\mathbf{X}\mathbf{Y}}^{X \text{ at } z}$  for various  $\kappa$  given by  $\kappa/N \in [1/8, 1/4, 1/2, 1]$  shown as  $\square$ ,  $\circ$ ,  $\diamond$ , and  $\triangle$  respectively.

a range of nearby locations according to  $\lambda$ . In essence, with this method one is convoluting the data with some point spread function to produce an estimate of the observable at all locations  $z$  based on the discrete list of measured locations for classified particles.

The term “non-parametric” is actually a bit of a misnomer, as the evaluation of either method requires specification of a parameter representing the bandwidth of the resolution filter. While the bandwidth in  $z$  for the kernel density prediction is constant for a given  $\lambda$ , for a given  $\kappa$  that for the nearest neighbor prediction is not, based as it is on the rank of the distances in  $z$  rather than their values. In the limit of infinite bandwidth, such that the  $z$  dependence disappears, both models give a prediction equal to that of the Bayesian estimate  $\langle p_m^X \rangle_{m|\mathbf{X}\mathbf{Y}} = J/(J+K)$ , as indicated by  $\triangle$  in the figures—the reason for the peak normalization of the kernel basis is so the kernel density estimates equal  $J$  or  $K$  for all  $z$  is this limit. The selection of the optimal value of the bandwidth parameter requires definition of some metric for its merit, introducing yet another source of subjectivity into the methodology. In contrast, the only arbitrary elements of the MaxEnt method are the limits of the prior, everything else following from repeated applications of the rules of probability theory to the state of knowledge specified at the outset, and even those are not truly arbitrary when one considers the physical nature of the apparatus and the objects.

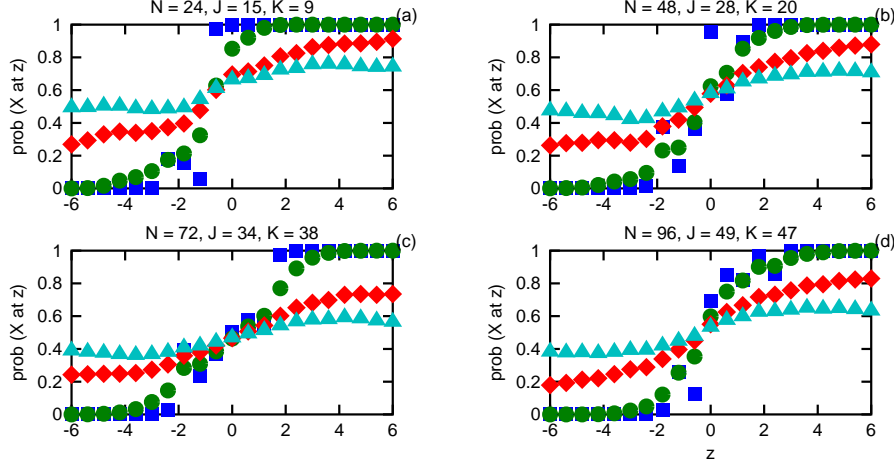


FIG 8. Comparison of the kernel density estimate predictions  $p_{z\lambda\mathbf{XY}}^{X\text{ at } z}$  for various  $\lambda$  given by  $\lambda \in [1/4, 1, 4, 16]$  shown as  $\square$ ,  $\circ$ ,  $\diamond$ , and  $\triangle$  respectively.

**6. Model mismatch.** What happens when the model used for the analysis of the data does not correspond in functional form to that of the underlying physical distribution? Without insight into the true nature of the objects measured, there is *a priori* no reason to suppose that some function chosen to resemble the data corresponds to that of the underlying physics. Specifically, let us consider a distribution for the location  $z < \bar{\phi}$  of particles of type  $X$  given by

$$(27) \quad f_{\Omega}(z_f | \tilde{\phi}) = z_f^{\tilde{\phi}-1} e^{-z_f/\tilde{\phi}} / \tilde{\phi}^{\tilde{\phi}} \Gamma(\tilde{\phi}),$$

where  $z_f \equiv \bar{\phi} - z$ , and similarly for  $g_{\Omega}(z_g | \tilde{\gamma})$  using  $z_g \equiv z - \bar{\gamma}$ , which one may recognize as a gamma distribution with a mean of  $\tilde{\phi}^2$  and a variance of  $\tilde{\phi}^3$  reflected in  $z$  and offset by  $\bar{\phi}$ . In panels (a) and (b) of Figure 9 are histograms of the locations of  $X$  and  $Y$  particles drawn from such a distribution with  $m_{\Omega} = 1$ ,  $\tilde{\phi}_{\Omega} = \tilde{\gamma}_{\Omega} = 1.5$ , and  $\bar{\phi}_{\Omega} = -\bar{\gamma}_{\Omega} = 4$ .

Using the model of the preceding section, we get the estimates for the observable  $\langle p_{z\mathbf{r}}^{X\text{ at } z} \rangle_{\mathbf{r}|\mathbf{XY}}$  and  $p_{z\mathbf{r}_0}^{X\text{ at } z} | \mathbf{XY}$  shown in panel (c) of Figure 9. While the observable resembles the physical distribution in the region of overlap between the particle types, the model is not designed to handle the case of finite boundaries in the location distribution. For the unidentified particles found outside that region, this model will almost always make the wrong prediction. One's suspicion might be aroused by noticing that the only measurements in the extreme regions are in contradiction to the model's



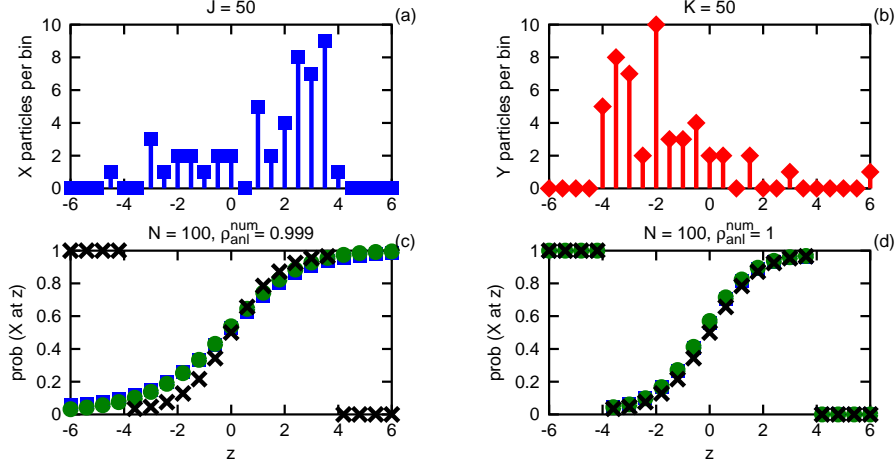


FIG 9. Histograms of location for type  $X$  in (a) and  $Y$  in (b) with a bin width of 0.5 for a set of 100 particles which yield the predictions  $\langle p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\mathbf{XY}}$  displayed as  $\square$  and  $p_{z\mathbf{r}_0}^{X \text{ at } z}|\mathbf{XY}$  displayed as  $\circ$  using the Gaussian distribution in (c) and the gamma distribution in (d) as well as the underlying distribution  $p_{z\Omega}^{X \text{ at } z}$  displayed as  $\times$ .

estimate. Of course, simply looking at the histograms reveals that the symmetric Gaussian model is not going to be the best fitting distribution for the identified particle locations.

Suppose now that somehow we gain knowledge of the functional form for  $f_{\Omega}(z)$  and  $g_{\Omega}(z)$ , for example by learning that the particles are racquetballs hit out of a tunnel such that  $X$  balls are bounced off the right wall and  $Y$  balls are bounced off the left, with the measurements for the  $z$  locations taken some distance away from the outlet of the tunnel. The parameters  $\tilde{\phi}_{\Omega}$  and  $\tilde{\gamma}_{\Omega}$  are assumed to be known from the geometry of the apparatus, with the parameters  $\mathbf{r} = (m, \tilde{\phi}, \tilde{\gamma})$  to be determined. The new parameter evidence is  $p_{\tilde{\phi}\mathbf{X}}^{\tilde{\phi}} \propto \tilde{\phi}^{-1} \prod_{z \in \mathbf{X}} f(z_f)$  when the prior  $p^{\tilde{\phi}} \propto \tilde{\phi}^{-1}$  is used, and similarly for  $\tilde{\gamma}$ . To remain finite as  $z_f \rightarrow 0$ , one requires  $\tilde{\phi} \geq 1$ , and the mode value  $\tilde{\phi}_0$  is easily found numerically, yielding the mode estimate  $p_{z\mathbf{r}_0}^{X \text{ at } z}$ , as shown in panel (d) of Figure 9. An upper limit of  $\tilde{\phi} \leq 4$  is consistent with the measurements, and the expectation of the observable  $\langle p_{z\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\mathbf{XY}}$  can be evaluated, also shown in panel (d). Having the correct physical model, even if its parameters are undetermined, is certainly an asset when attempting to use measurements to make predictions.

The point of this section is to emphasize how important one's knowledge of the situation is to the determination of one's results. The model one

selects should be based upon as much information as is available. When no single model presents itself as being physically correct, one must consider the alternatives, most often after having looked at the data—if the data look like an exponential decay, there is not much point in fitting a Gaussian. The prior for its parameters likewise should draw upon that background knowledge yet remain as unbiased as possible towards the final outcome. What Bayesian methods provide is a systematic framework for the comparison of models which all do a reasonable job of fitting the data, especially when the number or quality of measurements is low. The MaxEnt principle supplements Bayes' theorem by providing a systematic framework for the evaluation of the least biased prior based upon similarity transformations of the model with respect to the data.

**7. Unknown means and deviations with measurement uncertainty.** Returning to the use of the Gaussian model for the particle locations, let us now suppose the even more realistic situation where the location measurements are themselves subject to Gaussian deviation  $\sigma$ , presumed to be known from calibration of the measurement device, which for convenience will be set equal to the unit for  $z$  such that  $\sigma \equiv 1$ . The particle locations are supposed to be drawn from independent distributions as before, whose parameters are to be determined. The model selection ratio can be written in terms of the relative model likelihoods as

$$(28) \quad \rho_h^{fg} \equiv \frac{\langle p_{\bar{f}\bar{f}}^{\mathbf{X}} \rangle_{\bar{f}\bar{f}} \langle p_{\bar{g}\bar{g}}^{\mathbf{Y}} \rangle_{\bar{g}\bar{g}}}{\langle p_{\bar{h}\bar{h}}^{\mathbf{XY}} \rangle_{\bar{h}\bar{h}}} = \frac{\rho_f \rho_g}{\Delta_z \Delta_{\log \sigma} \rho_h} ,$$

with the prior limits notated as before. Focusing on the model for type  $X$ , each datum likelihood must now be expressed as an integral over all possible values of  $X_j$  according to the resolution of the apparatus,

$$(29a) \quad p_{\sigma \bar{f} \bar{f}}^{X_j} = \int_{-\infty}^{\infty} f(z_j) p_{\sigma X_j}^{z_j} dz_j$$

$$(29b) \quad = \int_{-\infty}^{\infty} (2\pi \tilde{f})^{-1} \exp^{-1/2} \{ [(\bar{f} - z_j)/\tilde{f}]^2 + (X_j - z_j)^2 \} dz_j$$

$$(29c) \quad = [2\pi(1 + \tilde{f}^2)]^{-1/2} \exp^{-1/2} [(\bar{f} - X_j)^2 / (1 + \tilde{f}^2)] .$$

The parameter evidence is now given by

$$(30) \quad p_{\sigma \mathbf{X}}^{\bar{f}\bar{f}} \propto \tilde{f}^{-1} [2\pi(1 + \tilde{f}^2)]^{-J/2} \exp^{-J/2} [(\bar{f}^2 - 2\bar{f}\langle X_j \rangle_j + \langle X_j^2 \rangle_j) / (1 + \tilde{f}^2)] ,$$

retaining explicitly the normalization of the likelihood but not the prior. The integral over  $\tilde{f}$  proceeds as before,

$$(31a) \quad p_{\sigma\mathbf{X}}^{\tilde{f}} = \int_{-\infty}^{\infty} p_{\sigma\mathbf{X}}^{\tilde{f}\tilde{f}} d\tilde{f}$$

$$(31b) \quad \propto \tilde{f}^{-1} [2\pi(1 + \tilde{f}^2)]^{(1-J)/2} J^{-1/2} \exp^{-1/2} [\xi_f^2 / (1 + \tilde{f}^2)] ,$$

yielding the marginal evidence for  $\tilde{f}$ . Under a change of variable  $\tilde{\sigma}^2 = 1 + \tilde{f}^2$  such that  $\tilde{f}^{-1} d\tilde{f} = \tilde{\sigma}(\tilde{\sigma}^2 - 1)^{-1} d\tilde{\sigma}$ , the marginal evidence for  $\tilde{f}$  can be rewritten as

$$(32) \quad p_{\sigma\mathbf{X}}^{\tilde{\sigma}} \propto (2\pi)^{(1-J)/2} J^{-1/2} (1 - \tilde{\sigma}^{-2})^{-1} \tilde{\sigma}^{-J} \exp^{-1/2} (\xi_f^2 / \tilde{\sigma}^2) ,$$

thus the remaining integral over  $\tilde{\sigma}$  can be written as an infinite series,

$$(33) \quad \rho_f \approx (2\pi)^{(1-J)/2} (8J)^{-1/2} \sum_{\alpha=0}^{\infty} 2^{(J+2\alpha)/2} \xi_f^{1-J-2\alpha} \Delta_{\Gamma}(\alpha, \xi_f) ,$$

where the approximation results from finite  $\Delta_z$  and  $\Delta_{\Gamma}$  now depends on  $\alpha$  as well as the limits of integration  $\tilde{\sigma}_{0,\infty} = (1 + \sigma_{0,\infty}^2)^{1/2}$ ,

$$(34) \quad \Delta_{\Gamma}(\alpha, \xi_f) \equiv \Gamma\left(\frac{J+2\alpha-1}{2}, \frac{\xi_f^2}{2\tilde{\sigma}_{\infty}^2}\right) - \Gamma\left(\frac{J+2\alpha-1}{2}, \frac{\xi_f^2}{2\tilde{\sigma}_0^2}\right) .$$

The relative likelihoods  $\rho_g$  and  $\rho_h$  are evaluated similarly.

The gradient of the parameter info,

$$(35) \quad \nabla q_{\sigma\mathbf{X}}^{\tilde{f}\tilde{f}} = \left[ \begin{array}{c} J(1 + \tilde{f}^2)^{-1} (\tilde{f} - \langle X_j \rangle_j) \\ \tilde{f}^{-1} + J(1 + \tilde{f}^2)^{-1} \tilde{f} [1 - (1 + \tilde{f}^2)^{-1} (\tilde{f}^2 - 2\tilde{f} \langle X_j \rangle_j + \langle X_j^2 \rangle_j)] \end{array} \right] ,$$

yields the same mode for the central location  $\tilde{f}_0 = \langle X_j \rangle_j$ . When that value is substituted into the expression for  $\partial_{\tilde{f}} q_{\sigma\mathbf{X}}^{\tilde{f}\tilde{f}}$ , the equation for the mode of the deviation becomes

$$(36) \quad 0 = (J+1)\tilde{f}_0^4 + (J+2-\xi)\tilde{f}_0^2 + 1 ,$$

whose root minimizes its contribution to the parameter info,

$$(37) \quad \tilde{f}_0 = \min_{\tilde{f}} \left[ \log \tilde{f} + \log(1 + \tilde{f}^2)^{J/2} + \xi_f^2 / 2(1 + \tilde{f}^2) \right] .$$

The remaining modes  $\tilde{g}_0$  and  $\tilde{h}_0$  are found similarly.

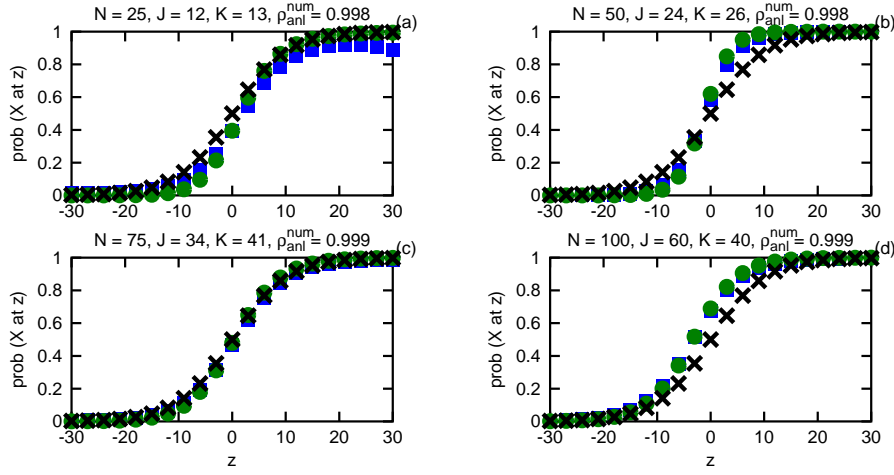


FIG 10. Comparison of the expectation value  $\langle p_{z\sigma\mathbf{r}}^{X \text{ at } z} \rangle_{\mathbf{r}|\sigma\mathbf{XY}}$  displayed as  $\square$  to the estimate from the parameter mode  $p_{z\sigma\mathbf{r}_0}^{X \text{ at } z}|\sigma\mathbf{XY}$  displayed as  $\circ$  as well as the underlying distribution  $p_{z\sigma\Omega}^{X \text{ at } z}$  displayed as  $\times$  with parameter values as given in Table 3.

The uncertainty in the measurement apparatus must also be taken into account when evaluating the observable. Given some measured location  $z$  for an unidentified particle, what we know is that its actual value  $z_\Omega$  is distributed around  $z$  with deviation  $\sigma$ . Consequently, the integration is now over not only the 5 dimensional parameter manifold but also over all possible values of  $z_\Omega$ ,

$$(38) \quad p_{z\sigma\mathbf{XY}}^{X \text{ at } z} = \int_{-\Delta_z/2}^{\Delta_z/2} p_{z_\Omega\sigma\mathbf{XY}}^{X \text{ at } z_\Omega} p_{z_\Omega}^{z_\Omega} dz_\Omega = \int_{-\Delta_z/2}^{\Delta_z/2} \int_{\mathbf{r}} p_{z_\Omega\mathbf{r}}^{X \text{ at } z_\Omega} p_{\sigma\mathbf{XY}}^{\mathbf{r}} p_{z_\Omega}^{z_\Omega} d\mathbf{r} dz_\Omega ,$$

which is evaluated numerically as before. The restricted limits of integration are found for the independent submanifolds such that the net integration measure is approximately normalized,  $\rho_{\text{anl}}^{\text{num}} \approx 1$ .

In Figure 10 we show the results of such an integration for  $N$  particles as indicated above each panel generated using parameters found in Table 3 with  $\Delta_z = 60$ . We can see that, despite the additional dimension of integration, the results are quite similar to what we had before. That should be no surprise, as the normal distribution has the feature that its mean and mode are at the same value. If during the calibration procedure a resolution (point spread) function other than Gaussian is determined for the apparatus, it should of course be used instead.

TABLE 3  
Parameters corresponding to Figure 10

panel	$\log_{10} \rho_h^{fg}$	$N$	$m_\Omega$	$\bar{f}_\Omega$	$\bar{g}_\Omega$	$\tilde{f}_\Omega$	$\tilde{g}_\Omega$
	$J$	$K$	$m_0$	$f_0$	$g_0$	$\tilde{f}_0$	$\tilde{g}_0$
a	1.99	25	1	10	-10	10	10
	12	13	0.923	10.7	-7.41	7.38	9.25
b	11.1	50	1	10	-10	10	10
	24	26	1	11.3	-14.1	7.56	8.02
c	8.98	75	1	10	-10	10	10
	34	41	1.14	8.98	-9.45	8.64	9.08
d	14.9	100	1	10	-10	10	10
	60	40	0.639	9.59	-12.5	9.55	9.52

**8. Reliability of the estimate.** So far we have danced around the topic of determining the reliability of the estimate strictly from the measurements at hand. In all the previous figures we have displayed the underlying physical distribution as a means of establishing that the method does indeed approach the “true” value as the number of measurements increases. However, in the real world, that knowledge is beyond our ken; we must make do with what the data have to say for themselves. That information is encoded in the evidence density for the model parameters given the measurements and the resolution of the apparatus.

When sufficient data exist that the significant evidence is restricted to some tiny region around a mode that barely moves as more data is collected, then we may as well call the prediction from the mode our single best estimate of the underlying distribution,  $p_{z\sigma\mathbf{r}_0}^{X\text{ at } z} \approx p_{z\sigma\Omega}^{X\text{ at } z}$ . In that case, the chance of making the correct prediction for some new datum is given by a simple truth table. Using the notation  $x_\Omega \equiv p_{z\sigma\Omega}^{X\text{ at } z}$ , we have

$$(39) \quad p(p_{z\sigma\mathbf{r}_0}^{X\text{ at } z} \text{ is correct}) \approx \text{Tr} \begin{bmatrix} x_\Omega^2 & x_\Omega(1-x_\Omega) \\ x_\Omega(1-x_\Omega) & (1-x_\Omega)^2 \end{bmatrix} = x_\Omega^2 + (1-x_\Omega)^2,$$

which indicates that even knowledge of the physical distribution does not guarantee a certain prediction for the particle type of the new datum; to be absolutely certain of the new particle’s type for any location, one must measure its classification. The error rate is greatest at the location where the particles appear with equal likelihood, as the chance of a successful prediction there is 1/2. When the mode of the model parameters is extremely well determined by the data, the error rate lies between 0 and 50% according to the value of  $p_{z\sigma\mathbf{r}_0}^{X\text{ at } z}$ .

While the prediction from the mode is the most likely contribution, the prediction from the mean of the observable is what encodes our best inference

about its value into a single number. If we want to know more about the observable, we have to do more work. As mentioned earlier, to describe the distribution around the expected value of the observable at some location, we need an additional parameter for its width and some function for its shape. The natural distribution for unit normalized positive quantities is the beta distribution  $p_{ab}^x = x^{a-1}(1-x)^{b-1}/\beta(a,b)$  that we have encountered in various guises already. Here, the problem is to determine its parameters  $a_z$  and  $b_z$  given empirical knowledge of the distribution of  $x_z \equiv p_{z\sigma}^{X \text{ at } z}$ .

One can easily show that the maximum likelihood estimate of the parameters is given by the solution of the system of equations

$$(40a) \quad \Lambda_1(a_z) - \Lambda_1(a_z + b_z) = \langle \log x_z \rangle ,$$

$$(40b) \quad \Lambda_1(b_z) - \Lambda_1(a_z + b_z) = \langle \log(1 - x_z) \rangle ,$$

using the notation  $\Lambda_k(r) \equiv (\partial_r)^k \log \Gamma(r)$  for the polygamma functions with integer order  $k$  and real argument  $r$ , which selects the parameters for the beta distribution whose values of  $\langle \log x_z \rangle$  and  $\langle \log(1 - x_z) \rangle$  equal those estimated from the empirical measurements. We set aside (for now) the question of whether a nonuniform prior  $p^{ab}$  should be incorporated on the grounds that a lot of effort will be put into converging the integrals such that the prior should have little effect.

The task now is to evaluate those expectation values conditioned on the given set of data. Instead of a single observable at each displayed location, there are now two to calculate, the first given by

$$(41) \quad \langle \log x_z \rangle = - \int_{-\Delta_z/2}^{\Delta_z/2} \int_{\mathbf{r}} \log[1 + m\zeta(z_\Omega)] p_{\sigma\mathbf{X}\mathbf{Y}}^{\mathbf{r}} p_{z\sigma}^{z_\Omega} d\mathbf{r} dz_\Omega ,$$

and the second given by

$$(42) \quad \langle \log(1 - x_z) \rangle = \langle \log x_z \rangle + \int_{-\Delta_z/2}^{\Delta_z/2} \int_{\mathbf{r}} \log[m\zeta(z_\Omega)] p_{\sigma\mathbf{X}\mathbf{Y}}^{\mathbf{r}} p_{z\sigma}^{z_\Omega} d\mathbf{r} dz_\Omega .$$

From those two estimates one finds the corresponding  $a_z$  and  $b_z$  according to Equations (40a) and (40b) above. The integral over  $z_\Omega$  in the second term can be expressed analytically when taken over infinite limits,

$$(43) \quad \int_{-\infty}^{\infty} \log[m\zeta(z_\Omega)] p_{z\sigma}^{z_\Omega} dz_\Omega = \log \left( \frac{m\bar{f}}{\bar{g}} \right) + \frac{(z - \bar{f})^2 + 1}{2\bar{f}^2} - \frac{(z - \bar{g})^2 + 1}{2\bar{g}^2} ,$$

thereby reducing the amount of effort required for its evaluation.

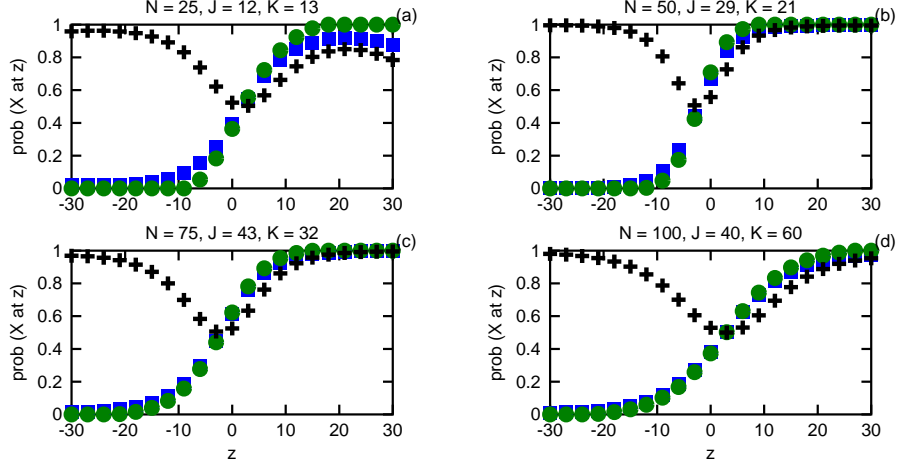


FIG 11. Comparison of the expectation value of the observable  $\langle x_z \rangle$  displayed as  $\square$  to its mode  $x_{z0}$  displayed as  $\circ$  as well as the expected chance that a prediction based on  $\langle x_z \rangle$  is correct displayed as  $+$  with parameter values as given in Table 4.

Having found  $a_z$  and  $b_z$ , the mean of the observable  $x_z$  may be determined according to  $\langle x_z \rangle = a_z / (a_z + b_z)$ , and its mode is given by  $x_{z0} = (a_z - 1) / (a_z + b_z - 2)$  when  $a_z, b_z \geq 1$  else may be found at 0 or 1. When using the expectation value to make predictions, the chance of a successful prediction is itself given by an expectation value,

$$(44) \quad \langle p(x_z \text{ is correct}) \rangle = \langle x_\Omega x_z \rangle + \langle (1 - x_\Omega)(1 - x_z) \rangle \approx \langle x_z \rangle^2 + (1 - \langle x_z \rangle)^2,$$

where the approximation results from using  $x_z$  as a proxy for  $x_\Omega$ . As promised, the Bayesian methodology has delivered an estimate of the observable used for predicting the type of unclassified particles as well as an estimate of its error rate based entirely upon the data at hand. In Figure 11 we show the mean and mode of the observable  $x_z$ , as well as the chance of a correct prediction according to Equation (44), for  $N$  particles as indicated above each panel generated using parameters found in Table 4.

**9. Discussion.** The purpose of the preceding exercises is to demonstrate how Bayes' theorem and the principle of maximum entropy are used to extract meaningful information from a set of data. When the observable representing the desired knowledge is not one of the parameters of the model, then the single best inference of its value is given by its expectation taken over the parameter manifold and weighted by the evidence at each parameter position. While the estimate from the parameter mode is the single

TABLE 4  
Parameters corresponding to Figure 11

panel	$\log_{10} \rho_h^{fg}$	$N$	$m_\Omega$	$\bar{f}_\Omega$	$\bar{g}_\Omega$	$\tilde{f}_\Omega$	$\tilde{g}_\Omega$
	$J$	$K$	$m_0$	$f_0$	$g_0$	$\tilde{f}_0$	$\tilde{g}_0$
a	1.99	25	1	10	-10	10	10
	12	13	0.923	10.7	-7.41	7.38	9.25
b	10.1	50	1	10	-10	10	10
	29	21	0.667	10.9	-13.7	7.95	8.17
c	10.2	75	1	10	-10	10	10
	43	32	0.705	9.58	-10.8	9.58	8.54
d	9.8	100	1	10	-10	10	10
	40	60	1.44	9.78	-8.71	10.0	10.3

most likely contribution, it is the mean of the observable which takes into account how much evidence there is in the data for other possibilities. When the MaxEnt principle is used to find the invariant measure for the parameter manifold, the results of a Bayesian analysis are the least biased possible, as each datum updates the evidence density according to its contribution to the available information which starts with the statement of the geometric properties of the model with respect to the data.

One feature of Bayesian data analysis as expressed in the language of conditional probabilities is that it forces one to specify the state of background knowledge upon which any estimate is based. Whenever real data are discussed, such statements necessarily include remarks on the nature of the measurement apparatus. That apparatus can be in the form of either a physical device such as a voltmeter or an abstract device such as a survey. While the parameter  $z$  has been given the interpretation of a spatial coordinate, it represents any observable that can be expressed in terms of a location-type parameter with uniform measure over an arbitrarily large domain. That parameter can of course be generalized to a data vector of any dimensionality according to the complexity of the situation at hand.

When one is not certain that the background knowledge includes specification of the physically correct functional form for the constituent distributions, there is a formal procedure for comparing the relative likelihood of competing models. That comparison is expressed not by the ratio of the peak likelihoods given by each model's parameter mode but rather by the ratio of the expected likelihoods given by the integral of the evidence over each model's parameter manifold. The distinction is important, as it is the latter which takes into account the principle of efficiency through the Occam factor. The most useful form of the parameter evidence density is that which drops the normalization of the prior yet retains any constant factors in the likelihood, as its integral appears as both the relative likelihood of the



model and the normalization constant for taking expectation values of observables. Care must be taken when dropping constants from an evaluation (whether numerical or analytic), as it is easy to make a mistake regarding the appropriate normalization for the given task.

When not all of the prior normalization factors cancel out of the model comparison ratio, then they must be accounted for explicitly. Quite often the MaxEnt prior will yield an infinite normalization when allowed to extend over an infinite domain, which is no less a problem for a maximum likelihood analysis that assigns any and all parameters a uniform prior. In that case, one is forced to consider more fully the nature of the apparatus as well as of the objects to be measured in order to determine sensible boundaries for one's prior state of knowledge. One generally hopes that finite boundaries do not truncate significantly the evidence density, but there can be times when physical constraints (such as positivity) dictate that what would otherwise be the mode lies outside the allowable parameter domain.

The process of inductive reasoning, through which one expresses one's degree of belief in the value of some observable, is not restricted to the simple distributions considered here but extends naturally to much more complicated situations. By going through the development of the methodology for increasingly less restrictive states of prior knowledge, we hope that readers have gained some insight into how to apply the method to problems they encounter. While almost everyone is comfortable with the process of finding the parameter mode, and most with taking expectation values according to the evidence density, the process of model selection often seems mysterious until one realizes that the relevant factors are just the relative likelihoods and any leftover prior normalizations.

**10. Conclusion.** In this article we have explored the MaxEnt approach to the problem of predicting the type of some unclassified particle given its location and a list of locations for particles of identified classification. The process of inductive reasoning is used to relate these quantities of interest, which incorporates both Bayes' theorem and the principle of maximum entropy to produce the least biased estimate from the data at hand. The expectation value of the observable is compared to that derived from the parameter mode, and they converge on the underlying distribution when sufficient data exists and the model function is known to be physically correct. When competing models must be compared, the procedure of evaluating the evidence ratio in terms of the relative likelihoods and prior normalizations has been explicitly determined.

The prediction from the Bayesian method is found to be superior to those

from non-parametric algorithms, such as nearest neighbor or kernel density estimation, especially when not much data exists for analysis. The reason is because conditional probability theory makes full use of all the available information rather than discarding the contribution from some datum on the grounds that it is far away from the desired location or smearing that information out according to some linear operator. The essential difference between inductive and deductive methods is that the former inverts the model to predict the data whereas the latter inverts the data to predict the model—see [Johnson \(2010\)](#) for an explicit example regarding sums of exponential functions and [Johnson \(2012\)](#) as regards the Fourier transform. Having invested in the presumably expensive classification of the particles in the training set, it makes sense that one should choose the most reliable method for analyzing that data. It is only those of us who generate data cheaply that worry about the expense of computing an integral.

A major advantage of the Bayesian method is that it is extensible. As new effects are brought into play, their parameters are included in the analysis in a straightforward manner. The language of conditional probability theory is sufficiently general that it can handle whatever state of background knowledge one specifies, including those which admit that competing models are available as well as those which admit that the measurements are themselves subject to error. “Reasoning in the face of uncertainty” is an apt description of inferential logic and is a problem often faced in the imperfect real world. The use of Bayes’ theorem requires one to specify the conditions upon which any and all probabilities are based, which seems like a lot of work at the outset but leads to a robust analysis upon completion. It also recognizes the difference between the most likely value of some observable and its expected value, which is important to consider when not many measurements are available. By working through the tasks of model selection, parameter estimation, and observable prediction explicitly for the case of these simple distributions, we hope the reader has learned how to apply the MaxEnt method to more complicated cases of data analysis.

## References.

- BERNTSEN, J., ESPELID, T. O. and GENZ, A. (1991). An adaptive algorithm for the approximate calculation of multiple integrals. *ACM Trans. Math. Softw.* **17** 437-451.
- BRETTTHORST, G. L. (1988). *Bayesian Spectrum Analysis and Parameter Estimation*. Springer-Verlag, Berlin, Germany.
- DOSE, V. (2003). Hyperplane priors. *AIP Conference Proceedings* **659** 350-360.
- DURRETT, R. (1994). *The Essentials of Probability*. Duxbury Press, A Division of Wadsworth, Inc., Belmont, CA.

- GENZ, A. C. and MALIK, A. A. (1980). Remarks on algorithm 006: An adaptive algorithm for numerical integration over an N-dimensional rectangular region. *Journal of Computational and Applied Mathematics* **6** 295-302.
- HALL, P., PARK, B. U. and SAMWORTH, R. J. (2008). Choice of neighbor order in nearest-neighbor classification. *Annals of Statistics* **36** 2135-2152.
- JAYNES, E. T. (1968). Prior probabilities. *IEEE Transactions On Systems Science and Cybernetics* **4** 227-241.
- JOHNSON, R. (2010). Fitting a sum of exponentials to lattice correlation functions using a non-uniform prior. *The European Physical Journal C - Particles and Fields* **70** 233-241.
- JOHNSON, R. (2012). MaxEnt power spectrum estimation using the Fourier transform for irregularly sampled data applied to a record of stellar luminosity. *Astrophysics and Space Science* **338** 35-48.
- PRESS, W., TEUKOLSKY, S., VETTERLING, W. and FLANNERY, B. (1992). *Numerical Recipes in C*, 2nd ed. Cambridge University Press, Cambridge, England.
- SIVIA, D. S. (1996). *Data Analysis: A Bayesian Tutorial*. Oxford Science Publications. Oxford University Press, Oxford, UK.

ROBERT W. JOHNSON  
ALPHAWAVE RESEARCH  
JONESBORO, GA 30238  
E-MAIL: [robjohnson@alphawaveresearch.com](mailto:robjohnson@alphawaveresearch.com)